Chapter 12

Incorrect Blood Pressure

Limbo: *In some Christian theologies, the eternal abode or state, neither heaven nor hell.*
  1. any intermediate, indeterminate state or condition.
  2. a place of confinement, neglect, or oblivion.

—*Webster's New World Dictionary,* Third College Edition, 1988

With fast highways and gloomy intersections behind us, we enter into the world of hospitals and operating rooms. The story here is about a medical device, a blood-pressure machine. These machines are very reliable and are used in wards and operating rooms. Nevertheless, the overall reliability of a system depends not just on whether the internal components work well. Reliability, as a design and evaluation criterion, should be extended to include correct interfaces and suitable user interaction.

The Blood-Pressure Machine

The device, as the name implies, automates the blood-pressure measurement sequence. Placing a cuff on the patient's arm is done manually by the nurse or physician, but the sequence of pumping the cuff, finding the peak pressure in the arteries, deflating the cuff, and measuring the lowest pressure—is all automatic. The device we will examine here is in the operating room. An operation will start in a short time, and while the patient is being prepared in an adjacent room, an anesthesiologist and his resident are checking the array of machinery scattered around. The boxy-looking blood-pressure machine sits on a metal cart by the surgery table (figure 12.1[a]).

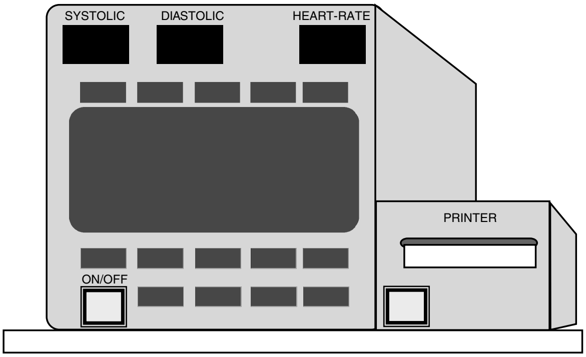

Figure 12.1 (a). The blood-pressure machine.

The machine comes to life by pressing the "on/off" button at the lower right of the interface (figure 12.1[a]). But then it takes about two minutes for the blood-pressure machine to go through an internal examination called a built-in test, which checks that all interior circuitry and software routines are intact. Such built-in tests are quite common in safety-critical systems. Their role is to check critical hardware and software functions, making sure the machine has no hidden failures, and then displaying the results of these internal checks in a short and binary message, either "BUILT-IN TEST—OK" or "BUILT-IN TEST—FAIL."

Today, at least for now, all is well. The built-in test is OK and the rectangular screen lights up with a glowing phosphorus green. Figure 12.1(b) shows that once the test is OK, the machine transitions automatically to OPERATIONAL.

Operational

Anesthesiologists and operating-room staff spend a considerable amount of time preparing and checking equipment to discover potential malfunctions, making sure that the equipment is operational and safe. In helping the surgeon and staff perform the operation safely, anesthesiologists have two goals: one is to induce and then maintain anesthesia throughout the surgery; the other is to detect, and promptly correct, deviations from the planned sequence of actions. Anesthesia, contrary to common belief, is not about "putting the patient to sleep." Rather it is a complex process that begins by gradually disarming the nervous system with extremely potent drugs; keeping the patient anesthetized, but alive, while he or she is subjected to excruciatingly painful surgical interventions; and then safely returning them to consciousness. Throughout this process, recording and assessing bodily signs such as blood pressure is vital.

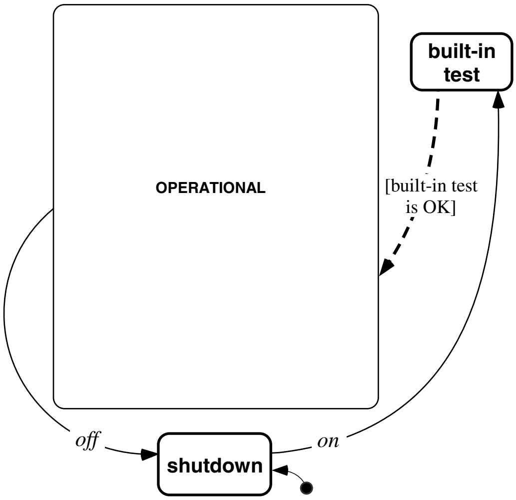

Figure 12.1 (b). From SHUTDOWN to OPERATIONAL mode (via the built-in test).

Now the patient is brought into the operating room and is quickly hooked up to the machines; the gray cuff is placed over his left arm, close to the heart.

Manual Measurement

The anesthesiology resident makes an initial measurement. He presses the "start" button (see figure 12.2[a]). The machine awakes and transitions into a SINGLE MEASUREMENT cycle (figure 12.2[b]). The cuff, which is now tightly wrapped around the patient's left arm, starts to inflate, producing a tightening pressure. Several seconds later, the machine has reached the maximum cuff pressure and stops. The cuff is fully inflated, and the brachial artery, under the triceps muscle, is blocked.

A hissing sound follows as the cuff deflates. The pressure on the arm is slowly relieved. In the old and familiar apparatus where the doctor pumps up the cuff with a black rubber bulb, the doctor listens carefully with the stethoscope, patiently waiting to detect the blood's return into the arm, which sounds like a turbulent stream as it passes by and overcomes the surrounding cuff's pressure. The heart, contracting and expanding, pumps blood into the arteries. It flows in a wave-like, undulating rhythm. The blood-pressure measures the *peak* and *trough* of that wave. In this blood pressure machine, a (Doppler) flow sensor is used to detect the blood's return. The cuff's pressure, at the point in which the blood starts returning, is equivalent to the *peak* pressure at which the heart pumps blood into the body. This peak pressure is the systolic pressure.

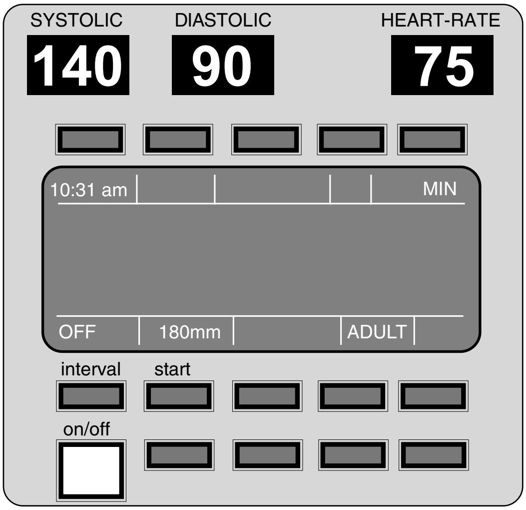

Figure 12.2 (a). Interface (in MANUAL mode).

As the pressure in the cuff is relieved, the brachial artery gradually opens up, and more and more blood gushes through. Seconds later, the cuff no longer restricts blood flow. At this point, the cuff's pressure is equivalent to the blood pressure inside the artery. This is the *resting*, or diastolic, pressure. The machine deflates a bit more, but the measurement is over. With the peak and rest values recorded and stored inside, the cycle is complete, and the machine transitions to OFF (figure 12.2[b]).

The systolic and diastolic pressures that correspond to the just completed measurement cycle are recorded and stored inside the machine and displayed on the top of the interface. In figure 12.2(a), the systolic blood pressure is at 140 millimeters of mercury and the diastolic is 90. The large rectangular screen, just below, provides information about the modes and setting of the machine: the current time, as you can see on the top-left corner, is 10:31 A.M., the patient is an ADULT, and the cuff is set to inflate all the way to 180 millimeters of mercury before it automatically deflates.

For the anesthesiologist, this is a wonderful labor-saving device, alleviating the time-consuming and attention-demanding measurement and recording with the old stethoscope and squeezer cuff. With the automated machine, all you do is press "start," and from there on everything is automatic. Furthermore, everyone in the operating room can look at the interface and see the patient's blood-pressure values.

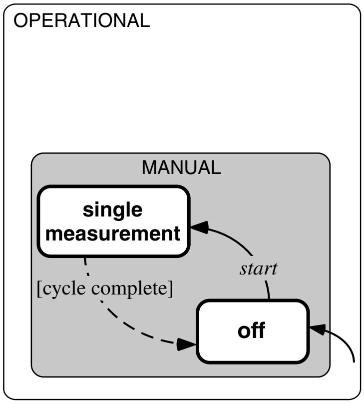

Figure 12.2 (b). Measurement cycle in MANUAL mode.

Automatic Measurement

Now the anesthesiologists are ready to sedate the patient with powerful drugs, bringing him slowly out of consciousness. During this descent, frequent measurements are required. So, short of pressing the "start" button every minute or two, the sequencing of measurements can also be done automatically, like an automatic timer. This way, the machine can run independently—initiating the measurement cycle, inflating and deflating—and then recording and presenting the blood-pressure measurement for display.

The anesthesiologist can choose among several automatic configurations: the machine can be set for CONTINUOUS operation, in which case the device is continuously measuring the patient's blood pressure (i.e., one cycle after another, no intervals). This setting is used if the patient's condition is severe and the surgeon and anesthesiologist need constant measurements. Another option is timer-interval, in which you set up the interval, or the hiatus, between successive measurements, and choose between 1, 2.5 , 5, 10, 15, 20, 30, 45, 60, or 120 minutes. Note that these intervals, as well as CONTINUOUS, are all reference values for the automatic mode.

Here is how the anesthesiologist actually sets up the machine for automatic mode to work: he first presses the "interval" button (lower left corner in figure 12.2[a]). This triggers an internal event, called "interval," which takes the machine from OFF to CONTINUOUS (see figure 12.3[a]). Another button press takes it to the one-minute interval setting. Press "interval" again and the machine is set for 2.5 minutes, another press and it's 5 minutes. The interval setting appears on the screen, just above the interval button (see figure 12.2[a]). During the course of the operation, the timer-interval is usually set for 5-, 10-, or 15-minute intervals, depending on the patient's condition. A

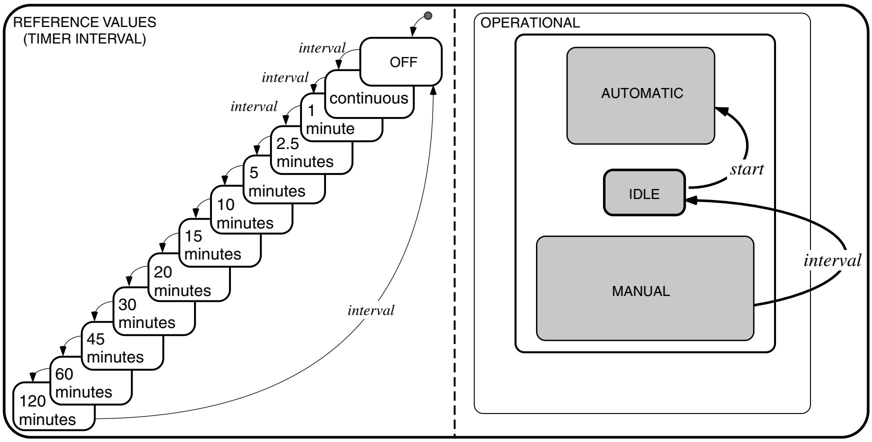

Figure 12.3. Modes and reference values.

common timer interval during the start of anesthesia is 5 minutes, and this was indeed the setting during the operation. (The longer interval settings, such as 60 minutes and 2 hours, are commonly used while the patient is in the recovery room.)

The event *interval*, in addition to setting the timer-interval value, also takes us out of MANUAL mode. As you can see in figure 12.3, the transition takes us initially into IDLE. And there the machine waits until the anesthesiologist presses the "start" button to begin the automatic measurement cycle.

At the completion of the first measurement cycle in automatic mode, the machine starts a time counter. (In figure 12.4[a], on the upper right corner of the screen, you can see the time counter—now showing that 2 minutes have elapsed.) Meanwhile, the anesthesiologist is carefully monitoring the patient's response to the drugs that he is providing through the intravenous tubes. By now the counter has reached five whole minutes. The blood-pressure machine wakes up and the cuff inflates with a sigh. At the end of this measurement cycle, everything looks normal—blood pressure is 135 over 83 and heart rate is at 70 (figure 12.4[a])—and now the counter starts again dutifully, counting up to the next measurement cycle.

Mode and Reference Values

The main task of the anesthesiologists is to assure the well-being of the patient through constant monitoring of the patient's vital signs. The process of anesthesia involves inducing sleep, bringing the patient down to a state of no pain (analgesia) and paralysis, and relaxation. Anesthesia is a very fragile place, kind of an in-between state. During anesthesia the intensity of pain is delicately

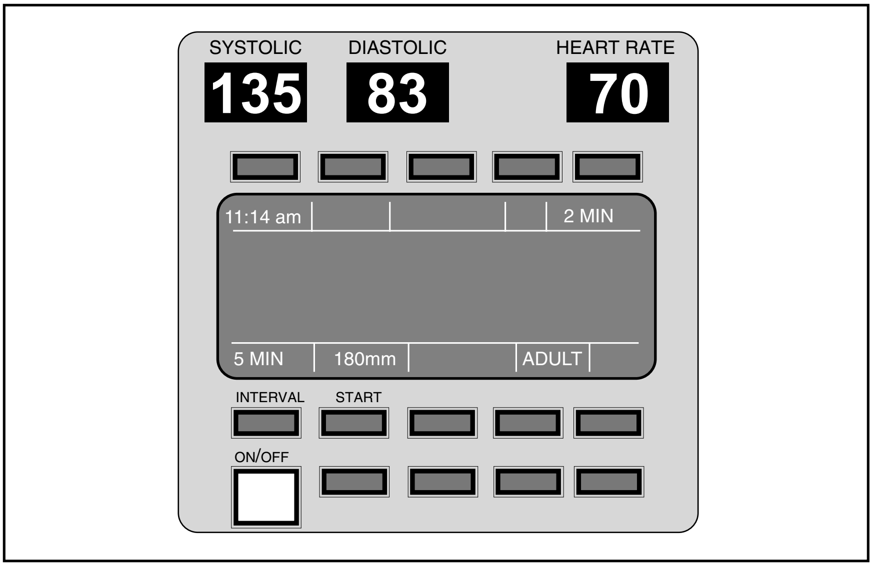

Figure 12.4 (a). In AUTOMATIC mode (2 minutes after start of measurement cycle).

countered by depth of sedation: if the patient is too awake he will experience excruciating pain, if he is too sedated he may be sliding down a dangerous and potentially lethal slope. Therefore, there is not a fixed amount of drugs that are used in anesthesia; rather, the drug quantities are titrated by their effects on the patient. Because of the close link between drug administration and the functioning of the heart, control of blood pressure is critical. And while the anesthesiologist and his resident are closely monitoring the patient's descent, and the machine cycles properly every 5 minutes, we can turn our attention back to the internal characteristics of the machine's behavior.

In considering user interaction with this device, we focus straight-away on the relationship between modes and reference values, which, as we have seen in previous chapters, is a source of many user-interaction problems. The event interval scrolls the machine between numerous interval reference values; now look at the event *interval* in figure 12.4(b). It appears there three times: it takes us out of manual mode into idle (*interval=continuous,* or 1 minute up to 120 minutes); it also transitions us from AUTOMATIC to IDLE, and finally, when we are in AUTOMATIC mode and would like to change back into manual engagement, we do this by pressing interval and setting it to off (*interval=off*).

The point is that the event *interval*—although appearing as an event that only changes (the timer interval) reference values—also switches modes. This multifunctionality of *interval* is a subtlety that will become crucial later. But for now, it is important to remember that quite frequently, during the operation, while the machine is in AUTOMATIC mode, the anesthesiologists may need to

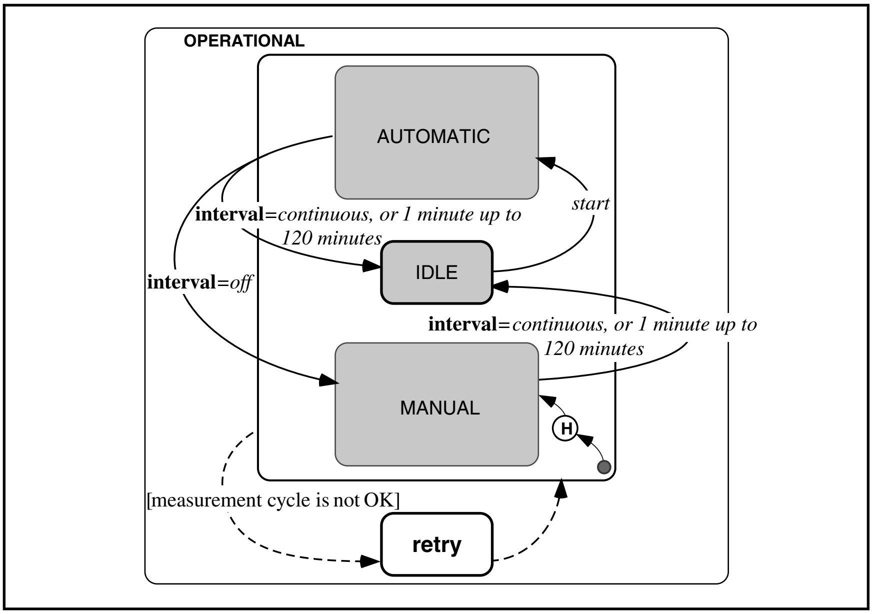

Figure 12.4 (b). Modes and envelope protection.

change the timer interval to either a higher or lower setting depending on the surgical needs and the status of the patient. To do that, they first need to select the desired timer-interval setting (say to 1 minute), and then press the "start" button to reactivate the machine in automatic mode (figure 12.4[b]).

Automatic Protection

Before we complete the description of the machine, it is worthwhile to mention that in addition to the presence of a watchful anesthesiologist, there is a built-in protection mechanism making sure that each measurement is performed properly and accurately. In figure 12.4(b), the model of the blood-pressure machine is expanded to include a protection mode. The RETRY mode is like a security blanket that covers both the manual and the automatic mode. While the machine is performing the cycle, there is a sensor and a software routine that constantly evaluate the quality of the measurement. Think of it as an ever-present quality-control team. If the measurement cycle is bad, it takes over and makes the machine repeat the cycle.

Such protection systems—that constantly make sure that the system does not wander beyond its zone of safe operation—are called *envelope protection systems*. You find them in machines that operate in safety-critical domains such as medicine, aviation, space flight, and nuclear power, making sure that these automated machines will work within the safe operating envelope and not produce bad outcomes. (In chapter 15, we will talk more about such envelope protection systems, in the context of aircraft and autopilots.)

Coming back to this blood-pressure machine, once the measurement triggered by the RETRY mode is okay, the machine automatically returns to normal operation. Figure 12.4(b) shows that the machine will return to the previously held mode. But how does the machine know which mode was previously active? The machine actually keeps track of this historical information (and we denote this ability to remember the previous mode with a history [H] transition that we have already discussed in detail in chapter 9). However, if the measurement triggered by the RETRY mode is still bad, the machine will try again. After three consecutive bad tries, the machine will abort automatically, and an "UNABLE TO MEASURE" message will appear on the screen.

A Critical Incident

Surgery can last for hours, while the anesthesiologists sit there and monitor the journey undertaken by the patient. But there is always a potential danger lurking just below the surface. Among bridge officers, anesthesiologists, pilots, and technicians who supervise critical automated systems, there is a colloquialism that the job is characterized by "hours of boredom—moments of terror." When those moments come, the anesthesiologist is usually faced with a complex situation that stems from the fact that many events and bodily functions are highly interconnected.

And things have gone *very wrong* with the surgery we have chosen to observe. The following is an excerpt from the investigation into a mishap that occurred with this specific blood-pressure machine:

"The patient's blood pressure read high ("hypertension") and the surgeon complained of profuse bleeding in the surgical field. The anesthesiologists set the device to read more frequently and became quite busy trying to lower the blood pressure to help the surgeon control the bleeding. Periodically the anesthesiologists glanced back at the blood-pressure device, only to see that it still showed an *elevated* blood pressure. They did not realize that it showed the same blood pressure as before and had never made another measurement. So, for 45 minutes they aggressively treated the hypertension with powerful drugs that dilated blood vessels. The bleeding continued to be a problem until the surgeon got control of the bleeding sites. Finally it was discovered that there had *not* been a recent measurement of blood-pressure, which was in fact found to be very *low*." (A low blood pressure is especially dangerous as vital organs may suffer irreversible damage if they are not sufficiently perfused by oxygenrich blood). In the end, the anesthesiologist was able to restore normal blood pressure and fortunately there was no lasting harm to the patient.

Analysis

Let's start by considering what happened here in the context of what we already know about this blood-pressure machine. At first, everything was normal and the machine was in a fully automatic mode at five-minute intervals. Later, the incident erupted: the blood pressure increased, and abdominal bleeding was out of control. The surgeon asked the anesthesiologists to reduce the patient's blood pressure, and they began administering drugs. But first, to better monitor the intervention, they set the machine to one-minute intervals.

What changed this episode from a routine intervention to an almost fatal complication was one (missing) button press. After resetting the timer interval from five minutes to one, the anesthesiologists probably forgot that a subsequent press on the "start" button was necessary to restart the measurement cycle.

At this point, we could place the blame on the anesthesiologists and classify it as "anesthesiologist error." But what assurance is there, for potential patients, that a similar error will not happen again? We need to go beyond the blame and try to find out what prompted this medical error, what was the sequence of events, and what can be learned from this incident that will help us better understand user interaction with automated systems.

Modes

From the user's perspective, the machine can operate in two distinct modes: (1) MANUAL, in which the user initiates each measurement cycle by pressing the "start" button, and (2) AUTOMATIC, in which the user pre-selects the interval and the machine initiates each measurement cycle automatically.

However, if you look carefully at figure 12.4(b), you will note that between MANUAL and AUTOMATIC there is another configuration. And this configuration is IDLE. When the machine is in MANUAL mode, selecting any interval setting sends the machine to this IDLE; a subsequent press on the "start" button is required to engage and activate the AUTOMATIC mode. And when the machine is already in AUTOMATIC, changing the interval (for example from five minutes to one), also sends the machine to IDLE, where it will stay, unless the user presses the "start" button. Therefore, for all practical purposes, we really have three different distinct modes here—MANUAL, AUTOMATIC, and IDLE. The user, however, is not fully aware of this intermediary IDLE mode. This, as you will see soon, played a significant role in the confusion that brought about this medical incident.

Feedback

Now let us consider the feedback provided to the user about the modes of the blood-pressure machine. When we are in MANUAL mode, the interval indication at the lower-left corner of the screen says OFF (see figure 12.2[a]); and when we are in AUTOMATIC, the same indicator shows the interval setting (five minutes in figure 12.4[a]). It appears straightforward and clear.

But consider what happens when we change the interval from five minutes to one. We press on the interval button, and now "1 MIN" is displayed on the screen. But has the machine transitioned to AUTOMATIC? Not really. Will the system sequence a future cycle? No; it will be stuck in IDLE (awaiting subsequent "start"). Nevertheless, the screen shows us the same indication *as if* we were in AUTOMATIC mode. The user assumes the machine is in the fully AUTOMATIC mode, when in fact the machine is stuck in a limbo mode—idling indefinitely in this in-between mode.

To see the problem, look at figure 12.5, where the indications (OFF, 1 MIN) are superimposed on top of the machine model. The resulting composite

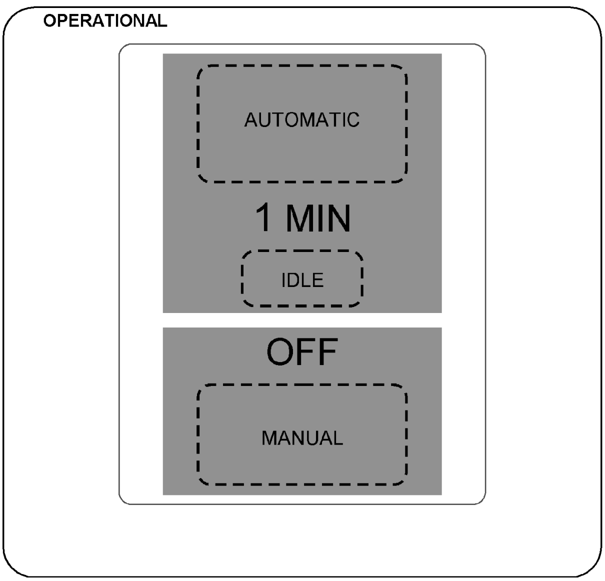

Figure 12.5. Superimposition of the interface indication on top of the machine; showing that the "1 MIN" indication covers both in AUTOMATIC and IDLE modes.

shows that the "1 MIN" indication can occur in both the AUTOMATIC and IDLE modes. But as we know well, in IDLE mode the machine will not make any measurement whatsoever. The interface provides the anesthesiologist with misleading information. Specifically, the interface is incorrect because two very distinct modes (AUTOMATIC and IDLE) are abstracted into a single indication ("1 MIN").

Population Stereotypes

Beyond the fact that the interface here is incorrect, it is also worthwhile to try to find out why the anesthesiologist or his resident forgot to press the "start" button immediately after selecting the one-minute interval. If these two user actions (select "interval" and press "start") had been conducted sequentially, the system would have gone immediately into AUTOMATIC mode. (This, of course, would not eliminate the incorrect interface problem—but it would certainly bypass it.)

Recall that all the anesthesiologists wanted to do was simply to change the timer-interval setting from five minutes to one. They were not after any mode changes. The kind of change they sought is not very different from changing the temperature setting in your toaster oven. For example, say that the salmon is not cooking fast enough and you want to increase the temperature. What do you do? On almost every oven, you simply grab the temperature knob and rotate it from 300 Fahrenheit to 450, and that's it. You are not expected to tell the system that you want it to stay in OVEN mode—you know it will.

As for light switches and TV/VCRs, there are strong population stereotypes when it comes to operating automated devices. But remember that there is nothing inherent in our brains about this or that population stereotype; the way we interact is simply based on our accumulated experience with automated devices. In particular, the relationship between a mode and its reference values, such that a simple change in reference value does not affect the mode, is one such prevailing stereotype. The anesthesiologists, like all of us, under stress and hurry, probably have fallen back into this stereotype. They assumed that changing the timer interval reference value is just that, and that the machine would stay in AUTOMATIC mode and measure at a faster rate. It didn't.

You see, the underlying design problem in this blood-pressure machine is the fact that the mode and reference-value settings are combined. The reference value is used both for setting the timer interval (such as five minutes and one) and for switching modes (MANUAL, AUTOMATIC, and OFF). Such designs usually result in hidden behaviors that many times even the designers of the system are unaware of. As a consequence, such designs are a breeding ground for confusion, frustration, and error. (In chapter 16 we will see a somewhat similar mode and reference value problem in autopilots of modern commercial aircraft that results in hidden behaviors.)

Recovery

There is a two-pronged approach to dealing with such mode and reference-value problems: one is to improve the design in order to stave off potential confusion. The other is to provide feedback that will allow for early detection and ways for a quick and safe recovery. So the next thing we will try to understand is why it took so long for the staff in the operating room to detect that the systolic/diastolic measurements were never updated, because with every minute that elapsed and drug dosages that were added, the patient's condition was getting worse. Just like the machine, he was also in limbo, holding tight, but slowly losing grip on that line between life and death.

Almost everyone in the room could see the blood-pressure values. The displays for the systolic and diastolic values, located above the screen, were rather large (about one inch high). However, there was no direct cue or salient feature anywhere on the interface to indicate the "age" of the values. Therefore, it is very likely that both the anesthesiologists and the rest of the surgical team believed that the blood-pressure measurement values were being updated, when in fact they were not.

The only indication that could cue the anesthesiologist that the device was stuck in idle mode was the small "time elapsed" indicator on the upper right corner of the screen (see figure 12.4[a]). This indicator displays the counter that starts ticking at the beginning of a new interval. Nevertheless, this was only an indirect cue that required that the anesthesiologist compare the elapsed time to the interval time. Apparently, it just wasn't powerful enough to compete with the prominent and compellingly misleading systolic and diastolic values on the interface.

In Closing

Several design features contributed to this critical incident: The root of the problem is the way modes and reference values were combined in this machine. As a result, it was necessary to design the interaction in such a way that required the user to always press the "start" button after changing the reference value, which, in turn, is a violation of a strong population stereotype. Finally, the lack of indication about the measurement's (systolic and diastolic) "age" hindered the surgical team's ability to recover from the error.

There are many potential solutions to the design problem in the blood-pressure machine. It may be possible to modify the design such that after changing the reference value, there is no need to press the "start" button. A more viable approach, one that both corresponds to the prevailing population stereotypes and provides salient mode feedback, would be to completely split the modes and reference values, such that the modes are clearly defined and displayed separately from the reference values. Finally, the indications of the systolic and diastolic values must include a feature that indicates the measurement's age, so that the nurses and doctors know when they are looking at old readings. It is noteworthy that the problems identified in this blood-pressure machine cannot be labeled as a manufacturing defect in the narrow sense of the term—because nothing broke and nothing worked different than advertised. Nevertheless, a design deficiency does exist in this machine. Furthermore, the resulting user interaction problem lies dormant, and when triggered, can easily turn a routine situation into a critical and harmful one.

There are numerous ways to improve the user interaction and the interface of the blood-pressure machine. Coming up with a "good" design is still an art form. But in terms of making sure that the design is good, we first want a correct interface. That is, it is essential that all the necessary information is there and that the interface does not tell you that the machine is in one mode when in fact it is in another. Second, we want to make sure that there are no violations of population stereotypes and there are no situations when, considering human interaction and human limitations, the interface "appears" nondeterministic and therefore incorrect to the user. Third, we need to make sure that there is consistency in the design of menus and that the labels on the buttons and interface are meaningful for the task. Finally, it is most important to conduct extensive usability testing with actual users and to evaluate the design in context. All of the above steps are important for consumer devices and are essential for safety-critical systems such as medical equipment.

Chapter 12

The incident described in this chapter took place several years ago at a large teaching hospital. The incident is detailed in a chapter, titled "Automation in anesthesiology," by David Gaba, M.D. which appeared in a book titled *Human Performance in Automated Systems: Current Research and Trends* by Mustapha Mouloua and Raja Parasuraman (Erlbaum, 1994, pages 57–63). The incident was not an isolated case, similar incidents have occurred in the past and the problem is known to many anesthesiologists. After the incident was investigated, the hospital management decided to remove this particular model of blood-pressure machine from all surgeries and wards.

In writing this chapter, I have relied on several academic publications on the topic of human factors in medicine. The book *Human Error in Medicine* (edited by Marilyn Bogner and published by Erlbaum, 1994) provided background information, as well as *Under the Mask: A Guide for Feeling Secure and Comfortable During Anesthesia and Surgery* by Dr. James Cottrell and Stephanie Golden (Rutgers University Press, 2001).

Chapter 13

The aircraft accident described in this chapter occurred several years ago. The factual information is based on the cockpit voice recorder and flight data recorder. The non-factual description and the painting of the scenes are based on my own flying experience. The actual evaluation of this procedure and the synchronization problem is somewhat more complicated than presented here, yet the results are the same. In writing this chapter I drew on previously published work conducted with my former advisor, Professor Earl Wiener, on the use and design of procedures. (See Asaf Degani and Earl Wiener, *The Human Factors of Flight-Deck Checklists: The Normal Checklist,* NASA Contractor Report number 177549, published in 1990; and *On the Design of Flight-Deck Procedures*, NASA Contractor Report number 177642, 1994.)

Problems in arming spoilers for landing have occurred in the past and contributed to many incidents and to a few accidents: In 1999, an American Airlines MD-80 aircraft crashed while landing on a wet and slippery runway in Little Rock, Arkansas. The American Airlines pilots also forgot to arm the spoilers before landing. Once they landed, as much as they tried, they were unable to stop the aircraft before the end of the runway. The aircraft overran the runway and broke in half, killing the captain and ten of his passengers. The full report on the American Airlines Flight 1420 accident can be obtained from the National Transportation Safety Board (NTSB) or downloaded from their web site (*American Airlines Flight 1420, Runway Overrun During Landing, Little Rock, Arkansas*, *June 1, 1999*; NTSB report number AAR-01/02).

Premature deployment of spoilers has also occurred in the past. On July 5, 1970, a McDonnell Douglas DC-8–63, operated by Air Canada, was making an approach to Toronto-Pearson International Airport. Sixty feet above the runway, the aircraft all of a sudden began to sink rapidly. The right outboard engine was torn off the aircraft in the subsequent heavy landing. The crew initiated a go-around and climbed to 3,000 feet. Then, a large piece of the right wing separated from the aircraft. The DC-8 stalled and crashed, killing all 109 people on board. The investigation report declared that the probable cause of the rapid descent while the aircraft was close to the runway was premature deployment of spoilers. As a result of this accident, the U.S. Federal Aviation Administration issued an Airworthiness Directive cautioning pilots against in-flight operation of ground spoilers by